

By chance is not enough: preserving relative density through non uniform sampling

Enrico Bertini, Giuseppe Santucci

Dipartimento di Informatica e Sistemistica - Università di Roma “La Sapienza”

Via Salaria, 113 - 00198 Roma, Italy - {bertini, santucci}@dis.uniroma1.it

Abstract

Dealing with visualizations containing large data set is a challenging issue and, in the field of Information Visualization, almost every visual technique reveals its drawback when visualizing large number of items. To deal with this problem we introduce a formal environment, modeling in a virtual space the image features we are interested in (e.g., absolute and relative density, clusters, etc.) and we define some metrics able to characterize the image decay. Such metrics drive our automatic techniques (i.e., not uniform sampling) rescuing the image features and making them visible to the user. In this paper we focus on 2D scatter-plots, devising a novel non uniform data sampling strategy able to preserve in an effective way relative densities.

Keywords—visual clutter, metrics, non-uniform sampling.

1 Introduction

Visualizing large data sets results, very often, in a cluttered image in which a lot of graphical elements overlap and many pixels become over plotted, hiding from the user the main image visual features.

We deal with this problem providing a formal framework to measure the amount of decay resulting from a given visualization, then we build, upon these measures, an automatic non uniform sampling strategy that aims at reducing such a degradation. We focus on a very common visual technique, 2D scatter-plots, analyzing the loss of information derived by overlapping pixels.

In this paper we improve and extend some preliminary results presented in [1], defining a formal model that estimates the amount of overlapping elements in a given area and the remaining free space. These pieces of information give an objective indication of what is eventually visualized on the physical device; exploiting such measures we can estimate the quality of the displayed graphic devising techniques able to recover the decayed visualization.

To eliminate the sense of clutter, we employ a low grain non uniform sampling technique dealing with the challenging issue of devising the *right* amount of sampling in order to preserve the visual characteristics of the underlying data. It is quite evident, in fact, that a too strong sampling is useless and destroys the less dense areas, while a too light sampling does not reduce the image clutter. The for-

mal model we discuss in the paper gives precise indications on the right amount of data sampling needed to produce a representation preserving the most important image characteristics, i.e., relative densities that are one of the main clues the user can grasp from 2D scatter-plots.

The contribution of this paper is twofold: (1) it presents a formal model that allows for defining and measuring data density both in terms of a virtual space and of a physical space (e.g., a display) and (2) it defines a novel automatic non uniform sampling technique driven by some metrics defined above the previous figures.

The paper is structured as follows: Section 2 analyzes related works, Section 3 describes the model we use to characterize clutter and density, formalizing the problem and introducing the metrics we are interested in, Section 4 describes our non uniform sampling technique, Section 5 discusses the results obtained applying our techniques to a real data set, and, finally, Section 6 presents some conclusions, open problems and future work.

2 Related Work

This paper deals with issues concerning metrics for Information Visualization and techniques to address the problem of overlapping pixels and visual clutter in computer displays. In the following we illustrate the research proposals closer to our approach and their relationship with our work.

2.1 Metrics for Information Visualization

As expressed in [6], Information Visualization needs metrics able to give precise indications on how effectively a visualization presents data and to measure its *goodness*. Some preliminary ideas have been proposed considering both formal measurements and guidelines to follow.

Tufte proposes in [9] some measures to estimate the quality of 2D representations of static data. Measures like the *lie factor*, that is the ratio of the size of an effect as shown graphically to its size in the data, are examples of first attempts to systematically provide indications about the quality of the image displayed. Tufte’s proposal however applies to paper based 2D visualizations and does not directly apply to interactive computer-based images. Brath in [7], starting from Tufte’s proposal, defines new metrics for static digital 3D images. He proposes metrics such as

data density (number of data points/number of pixels) that resemble Tufte’s approach together with new ones, aiming at measuring the visual image complexity. The *occlusion percentage*, for example, has connections with our work. It provides a measure of occluded elements in the visual space suggesting to reduce such a value as much as possible. These metrics are interesting and are more appropriate for describing digital representations. However, as stated by the author, they are still immature and need refinements.

While the above metrics aim at measuring a general goodness or at comparing different visual systems, our aim is to measure the accuracy of the visualization, that is, how well it represents the characteristics hidden inside data. They present some similarities with past metrics but operate at a lower level dealing with pixels and data points, providing measures that can directly be exploited to drive corrective actions. It is worth to note that, on the contrary of the above proposals, we will show how the suggested metrics can be exploited *in practice* to take quantitative decisions about corrective actions and enhance the current visualization.

2.2 Dealing with overlapping pixels and clutter

The problem of eliminating visual clutter and overlapping pixels to produce intelligible graphics has been addressed by many proposals.

Jittering as stated in [8], is a widely adopted technique that permits to make apparent pixels that naturally map into the same position into the screen. The idea is to slightly change the position of overlapping points in order to render them all visible. Similarly, space-filling pixel-based techniques [5] distribute data points along predefined curves to avoid overlapping pixels, shifting them to positions that are as close as possible to the original one.

Transparency is also an interesting technique to overcome occlusion and reduce clutter, both in 3D [12] and 2D [3] visualizations. However, when dealing with pixel-based visualizations it is not possible to convey transparency at the level of single pixels, though it is useless.

Constant density visualization [10][11] is an interesting technique to deal with clutter. Exploiting the idea of generalized fisheye views [4], it consists in giving more details to less dense areas and less details to denser areas, allowing the screen space to be optimally utilized and to reduce clutter. The problems with this approach are that it requires the user to interact with the system, the overall trend of data is generally lost, and some distortions are introduced.

Sampling is used in [2] to reduce the density of visual representation. As the authors state, if the sampling is made in random way, the distribution is preserved and though it is still possible to grasp some useful information about data correlation and distributions, permitting “to see the overall trends in the visualization but at a reduced den-




Figure 1: *Plotting mail parcels*

sity”. Even if interesting, this idea is not free of drawbacks. In particular, when the data present particular distributions, i.e., the data set has both very high and very low density areas, choosing the right amount of sampling is a challenging task. Depending on the amount of sampling two problems can arise: 1) If the sampling is too *strong* the areas in which the density is under a certain level become completely empty; 2) If the sampling is too *weak* the areas with higher densities will still look all the same (i.e., completely saturated) and consequently the density differences among them will be not perceived by the user. A first proposal in this direction is in [1], where an automatic uniform sampling technique is presented, able to compute the optimal sampling ratio w.r.t. some quality metrics.

Our approach differs from the above proposals for three main aspects:

- it provides a sound model for defining in both a virtual and physical space several metrics intended specifically for digital images;
- it provides, on the basis of the above figures, some *quantitative* information about the image decay;
- it exploits such numerical results for automatically computing *where*, *how*, and *how much* to sample preserving, as much as possible, a certain visual characteristic.

3 Modeling Visual Density and Clutter

In this section we present the formal framework that aims at modeling the clutter produced by over-plotting data. Some preliminary issues about the matter are in [1]; here we show a refinement of that results.

We consider a 2D space in which we plot elements by associating a pixel to each data element mapping two data

attributes on the spatial coordinates. As an example, Figure 1 shows about 160,000 mail parcels plotted on the X-Y plane according to their weight (X axis) and volume (Y axis). It is worth noting that, even if the number of plotted items is little, the area close to the origin is very crowded (usually parcels are very light and little), so a great number of collisions is present in that area: the most crowded area contains more than 50,000 (about 30 %) of the whole dataset compressed in less than 1 % of the whole screen.

Exploiting well known results coming from the calculus of probability, we derive a function that estimates the amount of colliding points and, as a consequence, the amount of free available space. More formally, two points are in collision when their projection is on the same physical pixel. In order to derive such a function, we imagine to *toss* n data points in a random way on a fixed area of p pixels. This assumption is quite reasonable if we conduct our analysis on *small* areas.

To construct such functions we use a probabilistic model based on the parameters just described, that here we summarize for the sake of clarity:

- n is the number of points we want to plot;
- p is the number of available pixels;
- k is the number of collisions;
- d is the number of free pixels.

The probability of having *exactly* k collisions plotting n points on an area of p pixels, $Pr(k, n, p)$, is given by the following function:

$$\begin{aligned} & \frac{PERM[(\binom{p}{n-k})\binom{n-k+k-1}{k}]}{p^n} && \text{if } n \leq p \text{ and } k \in [0, n-1] \\ & && \text{or } n > p \text{ and } k \in [n-p, n-1] \\ & 0 && \text{if } n > p \text{ and } k \in [0, n-p] \end{aligned}$$

The function is defined only for $k < n$, because it is impossible to have more collisions than plotted points. Moreover, it is easy to understand that in some cases the probability is equal to zero: if $n > p$, because of we are plotting more points than available pixels, we must necessarily have some collisions. For example, if we have an area of 8×8 pixels and we plot 66 points, we must necessarily have at least 2 collisions, so $Pr(0, 66, 64) = 0$ and $Pr(1, 66, 64) = 0$.

The basic idea of the formula is to calculate, given p pixels and n plotted points, the ratio between the number of possible cases showing *exactly* k collisions and the total number of possible configurations.

The latter is computed considering all the possible ways in which it is possible to choose n points among p pixels,

i.e., selecting n elements from a set of p elements allowing repetitions (dispositions with repetitions: p^n).

Calculating the # *config with exactly k collisions* is performed in three steps. First we calculate all the possible ways of selecting $n - k$ non colliding points from p pixels (combinations *without* repetitions: $\binom{p}{n-k}$). After that, for each of such combinations, we calculate all the possible ways of hitting k times one or more of the $n - k$ non colliding points in order to obtain exactly k collisions, that corresponds to selecting k elements from a set of $n - k$ elements *with* repetitions (combinations *with* repetitions: $\binom{n-k+k-1}{k}$). Finally, because of we are interested in all the possible dispositions, we need to count the permutations (PERM) of these combinations. Unfortunately, because of the variable number of duplicates (e.g., it is possible to have k collisions hitting $k+1$ times the same pixel p_i , or k times p_i and two times pixel p_j , or $k-1$ times p_i , two times pixel p_j , and two times pixel p_k and so on) we were not able to express such permutations by a close formula.

From the above expression we derived, through a C program, a series of functions (see Figure 5) showing the behavior of the observed area as the number of plotted points increases. More precisely, we compute the available free space d (Y axis, as percentage w.r.t. p), the mean of colliding elements k (Y axis, as percentage w.r.t. n) for any given number of plotted points n (X axis, as percentage w.r.t. p). For example, if we have an area of 64 pixels, the graph tell us that plotting 200% (128) of p points will produce an average of 56.7% (72.5) collisions. On the other hand, if we plot 128 points having 72.5 collisions we can compute the free pixels d , as $d = 64 - (128 - 72.5) = 8.5$ (13.3%).

The behavior of the functions is quite intuitive: as the number of plotted points n increases the percentage of collisions increases as well while the free space decreases; roughly speaking, we can say that over plotting four times the screen results in a totally saturated display (1.6% of free space).

Such functions can tell us how much we are saturating the space or, as a more complex possibility, the way in which the display is able to represent relative densities and how much to sample the data to guarantee a prefixed visualization quality. This result is exploited in the next section and we clarify it through an example. Assume that we are plotting n points on the area A_1 turning on p_1 pixels and $2n$ points on the area A_2 turning on p_2 pixels. In principle, the user should perceive area A_2 as containing more (i.e., twice as many) points as area A_1 . Because of collisions, $p_2 \leq 2p_1$ and as n increases the user initially loses the information that area A_2 contains twice as many points as A_1 and for greater values of n the user is not able to grasp any difference between A_1 and A_2 . As a numerical exam-




Figure 2: Colliding elements percentage

ple, if we plot 64 and 128 points on two 8×8 areas, the pixels turned on in the two areas will be 40.55 and 55.43, so the ratio of displayed pixels is only 1.36. In order to preserve the visual impression that area A_2 contains twice as many points as A_1 accepting a decay of 20 per cent we have to sample the data (64 and 128 points) as much as 50 per cent resulting in 32 and 64 points that, once plotted, turn on 25.32 and 40.55 pixels, i.e., a ratio of 1.6 (20 per cent of decay).

3.1 Data densities and represented density

The previous results give us a way to control and measure the number of colliding elements. Before introducing our optimization strategy, we need to clarify our scenario and to introduce new figures and definitions.

We assume the image is displayed on a rectangular area (measured in inches) and that small squares of area A divide the space in $m \times n$ sample areas (SA) where density is measured. Given a particular monitor, resolution and size affect the values used in calculations. In the following we assume that we are using a monitor of 1280x1024 pixels and size of 13"x10.5". Using these figures we have 1,310,720 pixels and if we choose SA of side $l = 0,08$ inch, the area is covered by 20.480 (128x160) sample areas whose dimension in pixel is 8×8 . We consider small areas because of it makes the uniform distribution assumption quite realistic.

For each $SA_{i,j}$, where $1 \leq i \leq m$ and $1 \leq j \leq n$, we calculate two different densities: *real data density* (or, shorter, data density) and *represented density*.

Data density is defined as $d_{i,j} = \frac{n_{i,j}}{A}$ where $n_{i,j}$ is the number of data points that fall into sample area $A_{i,j}$. For a given visualization, the set of data densities is finite and discrete. In fact, if we plot a number n of data elements into the display, each $SA_{i,j}$ assumes a value $d_{i,j}$ that is within the finite and discrete set of values: $0, \frac{1}{A}, \frac{2}{A}, \dots, \frac{n}{A}$. In general, for any given visualization, a subset of these

values will be really assumed by the sample areas. For each value we can compute the number of sample areas in which that value is present and an histogram showing the distribution of the various data densities can be computed. For example, if we plot 100 data points into an area of 10 sample areas, we could have the following configuration: 3 sample areas with 20 data points, 2 sample areas with 15 data points, 2 sample areas with 5 data points.

Represented density is defined as $rd_{i,j} = \frac{p_{i,j}}{A}$ where $p_{i,j}$ is the number of distinct active pixels that fall into $SA_{i,j}$. The number of different values that a sample area can assume is heavily dependent on the size of sample areas. If we adopt sample areas of size 8×8 pixels, as described before, the number of different not null represented densities is 64. Thus, we can represent at most 64 different represented density values. It is quite obvious that, because of collisions, $rd_{i,j} \leq d_{i,j}$.

Using the above definitions we devised an effective set of quality metrics whose complete discussion, however, is out of the scope of this paper (see [1] for a practical use of these quality metrics for uniform sampling strategies).

The above metrics, together with the statistical results give us the means to devise an automatic non uniform sampling technique described in the next section.

4 Non uniform sampling

In [1] a uniform sampling strategy has been presented, showing its ability in improving an image readability. Applying the same amount of sampling to the whole image is quite straightforward but presents several drawbacks. As an example, it is quite obvious that sampling areas presenting very low data density is useless and potentially dangerous. Moreover it is quite evident that the most important clues a user can grasp from 2D scatter-plots are **differences** in densities and our opinion is that a non uniform sampling can preserve in a more efficient way such differences.

The problem of representing relative densities is the one of creating an optimal mapping between the set of the actual data densities and the set of available represented densities. Each data density must be associated to one of the 64 (under the hypothesis of 8×8 sample areas) available represented densities. Any given visualization is one particular mapping. Consider the case in which a visualization is obtained by displaying a large data set. It likely corresponds to a mapping in which higher densities are all mapped onto few single represented densities, the ones in which quite all pixels are active (pane saturation). This is why in that areas relative densities can not be perceived: a large number of high data densities is mapped onto very close values. Our idea is to investigate how these mappings could be changed in order to present to the user more information about relative densities accepting, to a certain extent, some distortion.

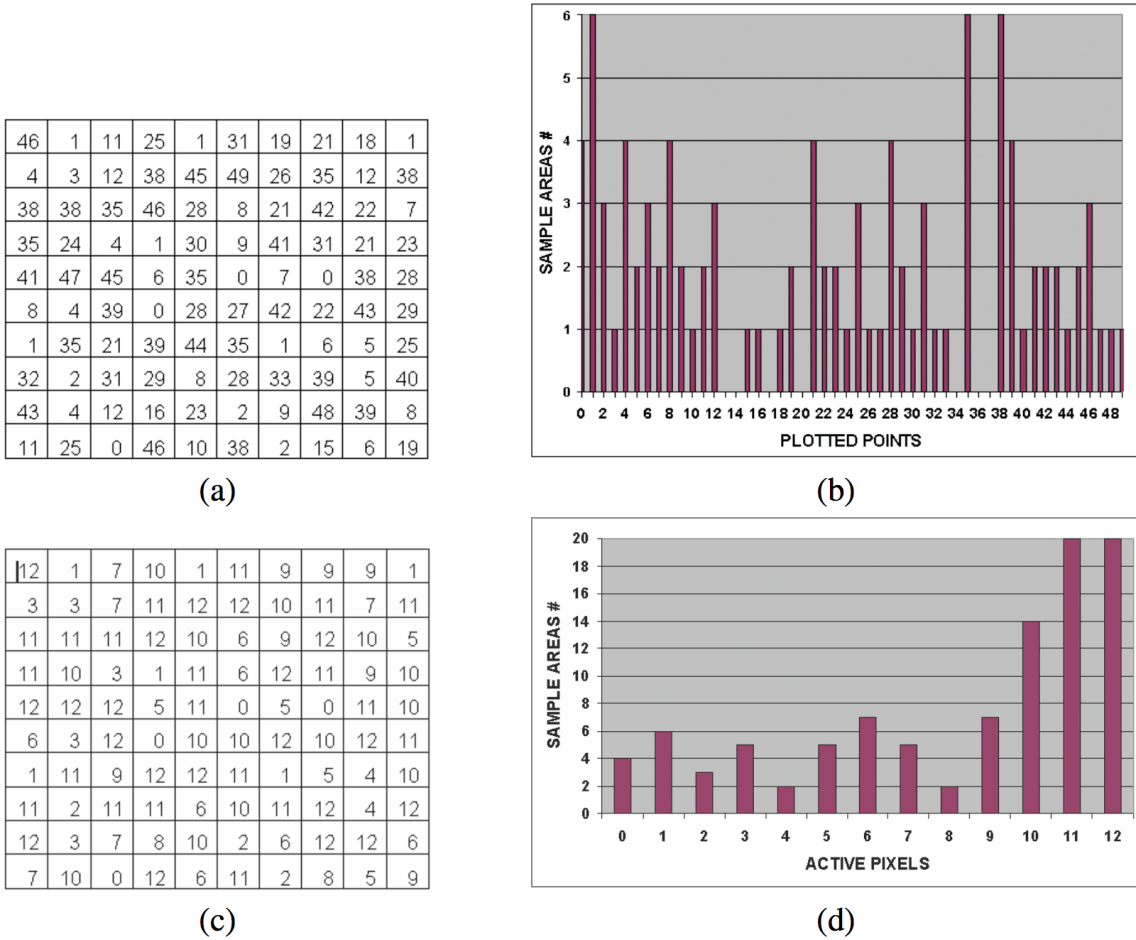


Figure 3: A screen area made of 100 sample areas: real (a) and represented (c) data densities

In the following we use a simple numeric example to clarify our approach. Assume we are plotting 2264 (this strange number comes from a random data generation) points on a screen composed by 400x400 pixels arranged in 100 sample areas of size 4x4 pixels. In the example we concentrate on the number of data elements or active pixels neglecting the SA area value (what we called A), that is just a constant. In Figure 3(a) the data densities (in terms of number of points) corresponding to each sample area are displayed.

Figure 3(b) shows the actual values of data densities (X axis) together with the associated number of sample areas sharing each value (Y axis). As an example, we can see that the maximum data density 49 is shared by just one sample area ($SA_{2,6}$) and the minimum data density 0 is shared by four sample areas ($SA_{5,6}$, $SA_{5,8}$, $SA_{6,4}$, $SA_{10,3}$). Figure 3 (c), obtained applying the statistical results discussed in Section 2 (see Figure 5), shows the actual represented density (in terms of active pixels) ranging, for


each $SA_{i,j}$, between 0 and 12. Looking at figure 3 (d) it is easy to discover that more than 50% of the visualization pane (54 sample areas out of 100) ranging between 22 and 49 data density collapsed on just three different represented data densities (10, 11, 12).

In order to improve such a situation we want to produce a new mapping among the given data densities and the 12 available represented densities. This can be done pursuing the goal of preserving the maximum number of differences, losing, on the other end, their extent. In other words, we want to present the user with as many difference in density as possible, partially hiding the real amount of such differences.

In order to obtain such a result, starting from figure Figure 3 (b) and considering only the 96 sample areas with data density > 0 we split the x axis in 12 (i.e., the available represented densities) adjacent *non uniform* intervals, each of them containing $96/12 = 8$ sample areas. Obviously, because of we are working on discrete values we cannot

12	1	4	6	1	8	5	5	5	1
2	2	4	10	12	12	7	9	4	10
10	10	9	12	7	3	5	11	6	3
9	6	2	1	8	4	11	8	5	6
11	12	12	3	9	0	3	0	10	7
3	2	11	0	7	7	11	6	12	7
1	9	5	11	12	9	1	2	2	6
8	1	8	7	3	7	8	11	2	11
12	2	4	5	6	1	4	12	11	3
4	6	0	12	4	10	1	5	3	5

(a)



(b)

Figure 4: *Non uniform sampling*

guarantee that each interval contains exactly 8 sample areas and we have to choose an approximation minimizing the variance. After that, the data elements belonging to the sample areas associated with the same interval i are sampled in a way that produces a represented density equal to i . As an example, the first interval encompasses data densities 1 (shared by 6 sample areas) and 2 (shared by 3 sample areas) and the associated data elements are sampled as much as needed in order to produce a represented density equal to 1. The second interval encompasses data density 3, 4, and 5 (7 sample areas) and, after the sampling, the resulting data density is 2, and so on.

The represented densities resulting from this approach are depicted in Figure 4 (a); Figure 4 (b) shows the new, more uniform distribution of such represented densities. We want to point out that in this new representation the above collapsed 54 data densities of Figure 3 (b) now range between 6 and 12 represented densities allowing the user to discover more density differences. On the other hand, as an example, the real difference between data densities 29 and 22 (1.32) is poorly mapped on represented densities 7 and 6 (1.16).

Roughly speaking, we can think at the whole process as follows. We have at disposal p different represented densities that are matched against k real data densities where, usually, $k \gg p$; that implies that each represented density is in charge to represent several, different data densities, hiding differences to the user. The game is to change, by non uniform sampling, the original data densities, altering their assignment to the p available represented densities in order to preserve the number of density differences.

5 Discussion

In this section we show the effectiveness of our technique commenting the images obtained applying different sampling strategies. We compare the images acquired visualizing a real dataset: the one containing 160,000 mail parcels already mentioned in Section 3.

The images come from a tool specifically developed for our purposes. It is a Java based application that permits to inspect several characteristics of the displayed image such as: the data/represented density of each sample area, some quality metrics, and the number of overlapping pixels. It is also possible to apply uniform and non-uniform sampling, and to filter sample areas with data/represented density out of a specific range.

Figure 5 shows: (a) the original visualization (no sampling), (b) the one obtained uniformly sampling the data leaving 80% of the original dataset, (c) the one obtained uniformly sampling the data leaving 20% of the original dataset (this value is the best uniform sampling ratio computed by the proposal shown in [1]), (d) the one obtained using non uniform sampling. It is quite evident that a too weak uniform sampling (Fig.5(b)) does not make apparent density differences in high density areas. Conversely, an optimized (but still too strong) uniform sampling (Fig.5(c)) makes them apparent but to the detriment of low density areas. In fact, the upper right area originally contained a cluster that is not visible anymore. Figure 5(d) shows the result obtained when applying non-uniform sampling. The features in the low density areas are still visible (as in the case of weak uniform sampling) but, at the same time, in the high density area it makes more evident density differ-




Figure 5: Comparison of various sampling methods visualizing the mail parcels dataset

ences that in the original image were not perceptible (as in the case of strong uniform sampling). Figure 5(e) makes it clearer. It is obtained filtering out the sample areas with data density lower than 810 (i.e., SA with less than 810 points) therefore showing the most dense areas. If compared with the other images it is easy to notice that while on Figure 5(d), that pattern is perfectly clear, on images (Fig. 5(a) and (b)) it is hidden in the saturated areas and on (Fig. 5(c)) it is faintly visible. Roughly speaking, we can say that our technique produce at the same time the advantages of both strong and weak sampling.

Another interesting aspect worth to mention, is how this technique can be operated. When applying uniform sampling, the choice of the amount of sampling to choose is critical. If the sampling factor is selected by hand, the user has to try many combinations until s/he finds the value that best conveys the information. To overcome this, we applied in [1] an algorithm to *automatically* devise the amount of sampling to apply. Exploiting the metrics presented in Section 3.1 we were able to find the best sampling factor to apply, but the problem of uniform sampling still held. Conversely, with non-uniform sampling there is no need to search into a space of solutions and the algorithm runs autonomously with the idea of assigning the available rep-

resented densities as *smartly* as it can.

The logic behind the algorithm can be better appreciated looking at Fig. 6 that compares the original and non uniformly sampled visualizations together with their densities histograms. The densities are more evenly distributed, allowing the dense areas to exhibit the underlying trends. Moreover, the peaks associated with the higher data densities (i.e., 62, 63, 64) are not present anymore.

6 Conclusions and future work

In this paper we presented a low grain, non uniform sampling sampling technique that automatically reduces visual clutter in a 2D scatter plot and preserves relative densities. To the best of our knowledge this approach is a quite novel way of sampling visual data. The technique exploits some statistical results and a formal model describing and measuring over plotting, screen occupation, and both data density and represented data density. Such a model allows for computing *where*, *how*, and *how much* to sample preserving some image characteristics (i.e., relative density).

Several open issues rise from this work:

- users must be involved. Our strategy provides precise figures but we need to map them against user




Figure 6: *Non uniform sampling*

perceptions. As an example, still referring to our approach, if a sample area contains twice as many active pixels as another one, does the user perceive the feeling of observing a double density for any total occupation of the areas? On the other hand, how much two sample areas may differ in pixel number still giving the user the sensation of having the same data density? We are currently designing some perceptive experiments, in order to deep this aspect. The next step will be to incorporate within our algorithms these issues.

- sampling areas. Several choices deserve more attention: it is our intention to analyze the influence of increasing/decreasing of sampling area dimension, in term of image quality and computational aspects.

We are actually extending the prototype functionalities to apply and verify our ideas. We want to implement a dataset generator to conduct controlled tests. The dataset generator will permit to generate artificial distributions, giving the possibility to control specific parameters, that will be used to create specific cases considered critical or interesting.

7 Acknowledgements

We would like to thank Pasquale Di Tucci for his invaluable help in implementing the software prototype.

References

- [1] E. Bertini and G.Santucci. Quality metrics for 2d scatterplot graphics: automatically reducing visual clutter. In *Proceedings of 4th International Symposium on SmartGraphics*, May 2004.
- [2] G. Ellis and A. Dix. Density control through random sampling: an architectural perspective. In *Proceedings of Conference on Information Visualisation*, pages 82–90, July 2002.
- [3] Jean-Daniel Fekete and Catherine Plaisant. Interactive information visualization of a million items. In *Proceedings of the IEEE Symposium on Information Visualization (InfoVis'02)*, page 117. IEEE Computer Society, 2002.
- [4] G. W. Furnas. Generalized fisheye views. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 16–23, 1986.
- [5] Daniel A. Keim and Annemarie Herrmann. The grid-fit algorithm: an efficient and effective approach to visualizing large amounts of spatial data. In *Proceedings of the conference on Visualization '98*, pages 181–188. IEEE Computer Society Press, 1998.
- [6] Nancy Miller, Beth Hetzler, Grant Nakamura, and Paul Whitney. The need for metrics in visual informa-

- tion analysis. In *Proceedings of the 1997 workshop on New paradigms in information visualization and manipulation*, pages 24–28. ACM Press, 1997.
- [7] Brath Richard. Concept demonstration: Metrics for effective information visualization. In *Proceedings For IEEE Symposium On Information Visualization*, pages 108–111. IEEE Service Center, Phoenix, AZ, 1997.
 - [8] Marjan Trutschl, Georges Grinstein, and Urska Cvek. Intelligently resolving point occlusion. In *Proceedings of the IEEE Symposium on Information Visualization 2003*, page 17. IEEE Computer Society, 2003.
 - [9] Edward R. Tufte. *The visual display of quantitative information*. Graphics Press, 1986.
 - [10] Allison Woodruff, James Landay, and Michael Stonebraker. Constant density visualizations of non-uniform distributions of data. In *Proceedings of the 11th annual ACM symposium on User interface software and technology*, pages 19–28. ACM Press, 1998.
 - [11] Allison Woodruff, James Landay, and Michael Stonebraker. Vida: (visual information density adjuster). In *CHI '99 extended abstracts on Human factors in computing systems*, pages 19–20. ACM Press, 1999.
 - [12] Shumin Zhai, William Buxton, and Paul Milgram. The partial-occlusion effect: utilizing semitransparency in 3d human-computer interaction. *ACM Trans. Comput.-Hum. Interact.*, 3(3):254–284, 1996.